

CLAIMS

What is claimed is:

1. A processor-implemented method for loading data from a vertical database table into a horizontal database table, comprising:

determining a number of rows in a vertical database table and a number of columns in a horizontal database table, wherein the vertical table includes at least three columns, with entries in a first column containing object identifiers, entries in a second column containing attribute names corresponding to the object identifiers, and entries in a third column containing attribute values corresponding to the attribute names, and the horizontal table includes a column for the object identifiers and attribute columns for the attributes names;

selecting one of a plurality of methods for reading data from the vertical database table and writing data to the horizontal database table based in part on the number of rows in the vertical database table and the number of columns in the horizontal database table; and

reading object identifiers and values of attributes from the vertical database table and writing the object identifiers and the values of attributes to the vertical database table using the selected one of the plurality of methods.

2. The method of claim 1, wherein access to the horizontal database table is provided by a database management system, the method further comprising:

in response to selection of a first one of the plurality of methods,

for each object identifier in the file,

generating a character string array having in character string format, data from the vertical database table including, an object identifier and attribute values for attributes associated with the object identifier; and

writing data from the character string array to a flat text file; and

loading data from the flat text file into the horizontal database using a utility function of the database management system, wherein the utility function is adapted to read data from a flat text file.

3. The method of claim 2, wherein generating a character string array includes determining a position for storage of an attribute value in the character string array using a hash of the attribute name associated with the attribute value.

4. The method of claim 2, wherein the first one of the plurality of methods is selected in response to the number of rows in the vertical table being greater than about 25,000 and the number of columns in the horizontal table being less than about 100.

5. The method of claim 1, wherein access to the horizontal database table is provided by a database management system, the method further comprising:

in response to selection of a first one of the plurality of methods,
for each object identifier in the file,
generating a character string array having in character string format, data from the vertical database table including, an object identifier and attribute values for attributes associated with the object identifier;
generating an SQL insert command from the contents of the character string array;
issuing the SQL insert command to the database management system; and
issuing an SQL commit command to the database management system.

6. The method of claim 5, wherein generating a character string array includes determining a position for storage of an attribute value in the character string array using a hash of the attribute name associated with the attribute value.

7. The method of claim 5, wherein the first one of the plurality of methods is selected in response to the number of rows in the vertical table being less than about 25,000 and the number of columns in the horizontal table being less than about 100.

8. The method of claim 1, wherein access to the vertical database table and access to the horizontal database table are provided by a database management system, the method further comprising:

in response to selection of a first one of the plurality of methods,
generating a single SQL command that selects data from the vertical table and inserts the data in the horizontal table for each object identifier in the vertical table and each column in the horizontal table; and
issuing the SQL command to the database management system.

9. The method of claim 8, further comprising:
 - determining an approximate percentage of a total number of attribute values that will have an undefined value in the horizontal table;
 - wherein the first one of the plurality of methods is selected in response to the approximate percentage being less than about 75%, the number of columns in the horizontal table being less than about 100, and the number of rows in the vertical table being less than about 5000.
10. The method of claim 9, further comprising:
 - in response to selection of a second one of the plurality of methods,
 - for each object identifier in the file,
 - generating a character string array having in character string format, data from the vertical database table including, an object identifier and attribute values for attributes associated with the object identifier; and
 - writing data from the character string array to a flat text file; and
 - loading data from the flat text file into the horizontal database using a utility function of the database management system, wherein the utility function is adapted to read data from a flat text file; and
 - wherein the second one of the plurality of methods is selected in response to the number of rows in the vertical table being greater than about 25,000 and the number of columns in the horizontal table being less than about 100; and
 - wherein the second one of the plurality of methods is selected in response to the number of columns in the horizontal table being less than about 100, the number of rows in the vertical table being in the range of about 3000 to 5000, and the approximate percentage being greater than about 75%.
11. The method of claim 1, wherein access to the vertical database table and access to the horizontal database table are provided by a database management system, the method further comprising:
 - in response to selection of a first one of the plurality of methods,
 - for each object identifier in the file,

generating a character string array having in character string format, data from the vertical database table including, an object identifier and attribute values for attributes associated with the object identifier; and

writing data from the character string array to a flat text file; and
loading data from the flat text file into the horizontal database using a utility function of the database management system, wherein the utility function is adapted to read data from a flat text file;

in response to selection of a second one of the plurality of methods,

for each object identifier in the file,

generating a character string array having in character string format, data from the vertical database table including, an object identifier and attribute values for attributes associated with the object identifier;

generating an SQL insert command from the contents of the character string array;

issuing the SQL insert command to the database management system; and
issuing an SQL commit command to the database management system;

in response to selection of a third first one of the plurality of methods,

generating a single SQL command that selects data from the vertical table and inserts the data in the horizontal table for each object identifier in the vertical table and each column in the horizontal table; and

issuing the SQL command to the database management system.

12. The method of claim 11, further comprising:

determining storage space requirements of the flat text file;

determining availability of sufficient retentive storage space based on the storage space requirements of for the flat text file;

selecting the first method in response to the number of columns being less than about 100, the number of rows being greater than about 5000, and there being available sufficient retentive storage space for the flat text file.

13. The method of claim 12, further comprising:

determining an approximate percentage of a total number of attribute values that will have an undefined value in the horizontal table;

in response insufficient retentive storage space being available for the flat text file,
selecting the second method in response to the number of rows being less than
about 8000 or the approximate percentage being greater than about 50%;
selecting the third method in response to the number of rows being greater than
or equal to about 8000 and the approximate percentage being less than or equal to about
50%.

14. An apparatus for loading data from a vertical database table into a horizontal database table, comprising:

means for determining a number of rows in a vertical database table and a number of columns in a horizontal database table, wherein the vertical table includes at least three columns, with entries in a first column containing object identifiers, entries in a second column containing attribute names corresponding to the object identifiers, and entries in a third column containing attribute values corresponding to the attribute names, and the horizontal table includes a column for the object identifiers and respective attribute columns for the attributes names;

means, responsive to the number of rows in the vertical database table and the number of columns in the horizontal database table, for selecting one of a plurality of approaches for reading data from the vertical database table and writing data to the horizontal database table; and

means for performing the selected one of the plurality of approaches.

15. The apparatus of claim 14, further comprising:

a database management system coupled to the horizontal database table;

means, responsive to a selection of a first one of the plurality of approaches, for generating, for each object identifier in the file, a character string array having in character string format, data from the vertical database table including, an object identifier and attribute values for attributes associated with the object identifier; and

means, responsive to selection of a first one of the plurality of approaches, for writing, for each object identifier in the file, data from the character string array to a flat text file; and

means, responsive to selection of a first one of the plurality of approaches, for loading data from the flat text file into the horizontal database using a utility function of the database management system, wherein the utility function is adapted to read data from a flat text file.

16. The apparatus of claim 14, further comprising:
 - a database management system coupled to the horizontal database table;
 - means, responsive to selection of a first one of the plurality of approaches, for generating, for each object identifier in the file, a character string array having in character string format, data from the vertical database table including, an object identifier and attribute values for attributes associated with the object identifier;
 - means for generating an SQL insert command from the contents of the character string array;
 - means for issuing the SQL insert command to the database management system; and
 - means for issuing an SQL commit command to the database management system.
17. The apparatus of claim 14, further comprising:
 - a database management system coupled to the vertical database table and to the horizontal database table;
 - means, responsive to selection of a first one of the plurality of approaches, for generating a single SQL command that selects data from the vertical table and inserts the data in the horizontal table for each object identifier in the vertical table and each column in the horizontal table; and
 - means for issuing the SQL command to the database management system.
18. A program storage medium, comprising:
 - at least one processor-readable program storage device configured with instructions for loading data from a vertical database table into a horizontal database table, wherein execution of the instructions by one or more processors causes the one or more processors to perform the operations including,
 - determining a number of rows in a vertical database table and a number of columns in a horizontal database table, wherein the vertical table includes at least three columns, with entries in a first column containing object identifiers, entries in a second column containing attribute names corresponding to the object identifiers, and entries in a third column containing attribute values corresponding to the attribute names, and the horizontal table includes a column for the object identifiers and respective attribute columns for the attributes names;

selecting one of a plurality of methods for reading data from the vertical database table and writing data to the horizontal database table based on the number of rows in the vertical database table and the number of columns in the horizontal database table; and

reading object identifiers and attributes values from the vertical database table and writing the object identifiers and attributes values to the vertical database table using the selected one of the plurality of methods.

19. The program storage medium of claim 18, wherein access to the horizontal database table is provided by a database management system, and the at least one processor-readable program storage device is further configured with instructions for execution by the one or more processors for performing the operations comprising:

in response to selection of a first one of the plurality of methods,

for each object identifier in the file,

generating a character string array having in character string format, data from the vertical database table including, an object identifier and attribute values for attributes associated with the object identifier; and

writing data from the character string array to a flat text file; and

loading data from the flat text file into the horizontal database using a utility function of the database management system, wherein the utility function is adapted to read data from a flat text file.

20. The program storage medium of claim 19, wherein the instructions for generating a character string array include instructions for determining a position for storage of an attribute value in the character string array using a hash of the attribute name associated with the attribute value.

21. The program storage medium of claim 19, wherein the first one of the plurality of methods is selected in response to the number of rows in the vertical table being greater than about 25,000 and the number of columns in the horizontal table being less than about 100.

22. The program storage medium of claim 18, wherein access to the horizontal database table is provided by a database management system, and the at least one processor-readable

program storage device is further configured with instructions for execution by the one or more processors for performing the operations comprising:

in response to selection of a first one of the plurality of methods,
for each object identifier in the file,
generating a character string array having in character string format, data from the vertical database table including, an object identifier and attribute values for attributes associated with the object identifier;
generating an SQL insert command from the contents of the character string array;
issuing the SQL insert command to the database management system; and
issuing an SQL commit command to the database management system.

23. The program storage medium of claim 22, wherein the instructions for generating a character string array include instructions for determining a position for storage of an attribute value in the character string array using a hash of the attribute name associated with the attribute value.

24. The program storage medium of claim 22, wherein the first one of the plurality of methods is selected in response to the number of rows in the vertical table being less than about 25,000 and the number of columns in the horizontal table being less than about 100.

25. The program storage medium of claim 18, wherein access to the vertical database table and access to the horizontal database table are provided by a database management system, and the at least one processor-readable program storage device is further configured with instructions for execution by the one or more processors for performing the operations comprising:

in response to selection of a first one of the plurality of methods,
generating a single SQL command that selects data from the vertical table and inserts the data in the horizontal table for each object identifier in the vertical table and each column in the horizontal table; and
issuing the SQL command to the database management system.

26. The program storage medium of claim 25, wherein the at least one processor-readable program storage device is further configured with instructions for execution by the one or more processors for performing the operations comprising:

determining an approximate percentage of a total number of attribute values that will have an undefined value in the horizontal table;

wherein the first one of the plurality of methods is selected in response to the approximate percentage being less than about 75%, the number of columns in the horizontal table being less than about 100, and the number of rows in the vertical table being less than about 5000.

27. The program storage medium of claim 26, wherein the at least one processor-readable program storage device is further configured with instructions for execution by the one or more processors for performing the operations comprising:

in response to selection of a second one of the plurality of methods,

for each object identifier in the file,

generating a character string array having in character string format, data from the vertical database table including, an object identifier and attribute values for attributes associated with the object identifier; and

writing data from the character string array to a flat text file; and

loading data from the flat text file into the horizontal database using a utility function of the database management system, wherein the utility function is adapted to read data from a flat text file; and

wherein the second one of the plurality of methods is selected in response to the number of rows in the vertical table being greater than about 25,000 and the number of columns in the horizontal table being less than about 100; and

wherein the second one of the plurality of methods is selected in response to the number of columns in the horizontal table being less than about 100, the number of rows in the vertical table being in the range of about 3000 to 5000, and the approximate percentage being greater than about 75%.

28. The program storage medium of claim 18, wherein access to the vertical database table and access to the horizontal database table are provided by a database management system, and

the at least one processor-readable program storage device is further configured with instructions for execution by the one or more processors for performing the operations comprising:

in response to selection of a first one of the plurality of methods,
for each object identifier in the file,
generating a character string array having in character string format, data from the vertical database table including, an object identifier and attribute values for attributes associated with the object identifier; and
writing data from the character string array to a flat text file; and
loading data from the flat text file into the horizontal database using a utility function of the database management system, wherein the utility function is adapted to read data from a flat text file;
in response to selection of a second one of the plurality of methods,
for each object identifier in the file,
generating a character string array having in character string format, data from the vertical database table including, an object identifier and attribute values for attributes associated with the object identifier;
generating an SQL insert command from the contents of the character string array;
issuing the SQL insert command to the database management system; and
issuing an SQL commit command to the database management system;
in response to selection of a third first one of the plurality of methods,
generating a single SQL command that selects data from the vertical table and inserts the data in the horizontal table for each object identifier in the vertical table and each column in the horizontal table; and
issuing the SQL command to the database management system.

29. The program storage medium of claim 28, wherein the at least one processor-readable program storage device is further configured with instructions for execution by the one or more processors for performing the operations comprising:

determining storage space requirements of the flat text file;
determining availability of sufficient retentive storage space based on the storage space requirements of for the flat text file;

selecting the first method in response to the number of columns being less than about 100, the number of rows being greater than about 5000, and there being available sufficient retentive storage space for the flat text file.

30. The program storage medium of claim 29, wherein the at least one processor-readable program storage device is further configured with instructions for execution by the one or more processors for performing the operations comprising:

 determining an approximate percentage of a total number of attribute values that will have an undefined value in the horizontal table;

 in response insufficient retentive storage space being available for the flat text file,

 selecting the second method in response to the number of rows being less than about 8000 or the approximate percentage being greater than about 50%;

 selecting the third method in response to the number of rows being greater than or equal to about 8000 and the approximate percentage being less than or equal to about 50%.